arXiv:1603.01150v1 [math.GR] 3 Mar 2016

THE BASILICA THOMPSON GROUP IS NOT FINITELY PRESENTED

STEFAN WITZEL AND MATTHEW C. B. ZAREMSKY

Abstract. We show that the Basilica Thompson group introduced by Belk and Forrest is not finitely presented, and in fact is not of type FP_2. The proof involves developing techniques for proving non-simple connectedness of certain subcomplexes of CAT(0) cube complexes.

Introduction

J. Belk and B. Forrest [BF15a] introduced the *Basilica Thompson group* T_B (defined in Definition 1.10 below). They showed that it is virtually simple, generated by four elements, and is a sub- as well as a supergroup of Thompson's group T. The question of whether it is finitely presented, however, remained open. In this paper we prove:

Theorem. *T_B is not finitely presented.*

The Basilica Thompson group is an example of a *rearrangement group of a fractal* as defined by Belk and Forrest in [BF15b]. These groups arise from *edge replacement systems* and act naturally on self-similar spaces, for example Julia sets. They generalize *Thompson's groups* F, T and V. For certain rearrangement groups Belk and Forrest proved that they are of *type* F_∞, confirming the "expected" behavior for relatives of Thompson's groups. As for the classical Thompson's groups, the proof relies on studying the action on an associated CAT(0) cube complex. The reason that (non-)finite presentability of T_B remained open is that the known local methods are not suited to prove the kind of negative connectivity statement needed. Our proof of the theorem involves a global analysis.

We now know of (virtually) simple rearrangement groups with two extremal finiteness properties: of type F_1 but not of type F_2, and of type F_∞. If the intermediate finiteness properties could also be achieved, e.g. of type F_n but not F_{n+1} for any n, this would provide an infinite family of pairwise non-quasi-isometric simple groups (the only other known examples of this kind are hyperbolic Kac–Moody groups [CR09]). We therefore ask:

Question. *Is there a rearrangement group that is virtually simple, finitely presented, but not of type F_∞?*

Date: March 4, 2016.

2010 *Mathematics Subject Classification.* Primary 20F65; Secondary 57M07.

Key words and phrases. Basilica Julia set, Thompson group, finite presentation, CAT(0) cube complex.

Acknowledgments. We are grateful to Jim Belk for posing this problem to us, to Belk and Brad Forrest for helpful discussions, and to the Binghamton University math department for its hospitality while the first named author was visiting the second named author. The first named author also gratefully acknowledges support by the DFG through the project WI 4079/2 and through the SFB 701.

1. Background

In this section we recall the necessary background material from Sections 1, 2 and 3 of [BF15b] on *edge replacement systems*, and the resulting groups and complexes. At the end we also recall some background on discrete Morse theory, and on CAT(0) cube complexes.

1.1. Edge replacement systems.

We will consider finite directed graphs G. The notation $V(G)$ and $E(G)$ will always mean the vertex and edges sets of G.

Definition 1.1 (Edge replacement rule)**.** An *(edge) replacement rule* is a pair $e \to R$, for e a non-loop directed edge, say with initial vertex v and terminal vertex w, and R a finite directed graph with $v, w \in V(R)$.

Definition 1.2 (Edge replacement system)**.** An *(edge) replacement system* is a pair $(G_0, e \to R)$, for G_0 a finite directed graph, called the *base graph*, and $e \to R$ a replacement rule. We will always assume $(G_0, e \to R)$ is *expanding* (see [BF15b, Definition 1.8]), i.e. G_0 has no isolated vertices, v and w do not share an edge in R, and $|V(R)| \geq 3$ and $|E(R)| \geq 2$.

Applying a replacement rule to the edge $\varepsilon \in E(G)$ of a graph G amounts to removing the edge ε and replacing it by R. Note that we do allow ε to be a loop. We denote by $G \lhd \varepsilon$ the graph obtained from G by replacing the edge ε by R in this way. We call $G \lhd \varepsilon$ a *simple expansion* of G. Any graph obtained from G by a finite sequence of simple expansions is called an *expansion* of G. The reverse of a (simple) expansion is called a *(simple) contraction*. We will "address" edges and vertices of expansions by concatenating addresses of edges and vertices from G and R; for example, if the edges of R are called $1, 2, 3$ then the new edges in the expansion $G \lhd \varepsilon$ are called $\varepsilon 1, \varepsilon 2, \varepsilon 3$, and the old edges retain their addresses from G. See Example 1.9 for pictures.

Definition 1.3 (Limit space)**.** Let $\mathcal{R} = (G_0, e \to R)$ be an edge replacement system. For each $n \in \mathbb{N}$ let G_n be the result of applying the replacement rule to each edge of G_{n-1} once. Let $\Omega := E(G_0) \times E(R)^\infty$ be the set of sequences of edges, with leading edge from G_0 and all others from R. Declare that two such sequences $\varepsilon_0 \varepsilon_1 \cdots$ and $\varepsilon_0' \varepsilon_1' \cdots$ are *equivalent* if for all n the edges of G_n with "addresses" $\varepsilon_0 \cdots \varepsilon_n$ and $\varepsilon_0' \cdots \varepsilon_n'$ share a vertex. Define the *limit space* X for $\mathcal{R}$ to be the space of equivalence classes $[\varepsilon_0 \varepsilon_1 \cdots]$ of such sequences.

The limit space X of a replacement system $\mathcal{R}$ is compact and metrizable ([BF15b, Theorem 1.24]). Our groups of interest are certain groups of homeomorphisms of such limit spaces, which is the subject of the next subsection.

1.2. **Rearrangement groups.** We now define *rearrangement groups*, which are certain groups of homeomorphisms of limit spaces.

Definition 1.4 (Cell)**.** Let X be the limit space of a replacement system $\mathcal{R} = (G_0, e \to R)$. Let $\varepsilon = \varepsilon_0 \cdots \varepsilon_n$ be an edge of some expansion of G_0. The *cell* $C(\varepsilon) \subseteq X$ is the subspace consisting of equivalence classes of sequences representable by a sequence with $\varepsilon_0 \cdots \varepsilon_n$ as a prefix.

Definition 1.5 (Canonical homeomorphism)**.** Let $C(\varepsilon)$ and $C(\varepsilon')$ be two cells such that ε and ε' are either both loops or both non-loops. The *canonical homeomorphism* $\Phi\colon C(\varepsilon) \to C(\varepsilon')$ is the map defined via the prefix replacement rule

$$\Phi([\varepsilon\zeta_1\zeta_2\cdots]) := [\varepsilon'\zeta_1\zeta_2\cdots].$$

Definition 1.6 (Rearrangement)**.** Let X be the limit space of a replacement system $\mathcal{R} = (G_0, e \to R)$. A homeomorphism $f\colon X \to X$ is called a *rearrangement* if there exist finitely many cells $C(\varepsilon_1), \ldots, C(\varepsilon_n)$ such that the cells cover X, have pairwise disjoint *interiors* (defined in [BF15b, Section 1.3]), and such that each restriction $f|_{C(\varepsilon_i)}$ is a canonical homeomorphism.

Definition 1.7 (Rearrangement group)**.** The rearrangements of X form a group [BF15b, Proposition 1.15], called the *rearrangement group* Γ of X.

Definition 1.8 (Graph pair diagram)**.** Let $f\colon X \to X$ be a rearrangement. A *graph pair diagram* for f is a triple (E_-, E_+, φ), where $E_\pm$ are expansions of G_0 and $\varphi\colon E_- \to E_+$ is a graph isomorphism, such that for every edge ε of E_- the restriction of f to $C(\varepsilon)$ is a canonical homeomorphism from $C(\varepsilon)$ to $C(\varphi(\varepsilon))$. The idea is that, even though f is a homeomorphism of X, it can already be realized at some finite expansion stage. In a graph pair diagram (E_-, E_+, φ), we call E_- the *domain graph* and E_+ the *range graph*.

Example 1.9 (The Basilica rewriting system)**.** Consider the replacement rule

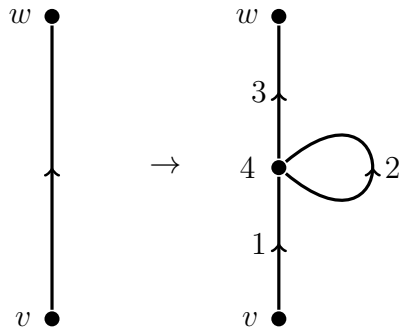

where $E(R) = \{1, 2, 3\}$ and $V(R) = \{4\}$. If G_0 is the graph

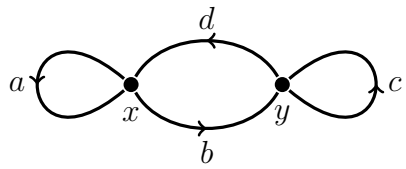

then for example the graphs

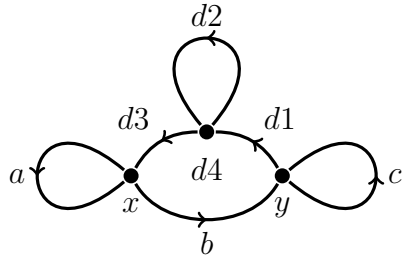

and

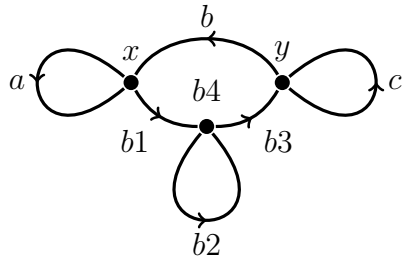

are expansions of G_0, with the addresses given for the edges and vertices. For X the limit space of the rearrangement system, the map $X \to X$ given piecewise by the canonical homeomorphisms

$$\begin{aligned} a* &\mapsto a* \\ b* &\mapsto b1* \\ c* &\mapsto b2* \\ d1* &\mapsto b3* \\ d2* &\mapsto c* \\ d3* &\mapsto d* \end{aligned}$$

is a rearrangement. The graph pair diagram for this map consists of the above graphs, with the isomorphism φ given by erasing the "$*$"'s in these canonical homeomorphisms.

Definition 1.10. The *Basilica Thompson group* T_B is the rearrangement group of the rewriting system in Example 1.9.

1.3. **Cube complexes for rearrangement groups.** In this subsection, we recall the CAT(0) cube complex on which a rearrangement group Γ acts. The replacement rule $e \to R$ will be fixed throughout. Given a base graph G_0, we will denote the cube complex by $K(G_0, e \to R)$. In [BF15b] it was denoted $K(\Gamma)$, but for our purposes it is important to keep track of the base graph G_0 used, and less important to keep track of the group Γ. For any choice of base graph G_0, denote by $X(G_0)$ the limit space of the replacement system $(G_0, e \to R)$. We extend the definition of rearrangement from Definition 1.6 as follows:

Definition 1.11 (Rearrangement)**.** Let G_0 and G be two choices of base graph, so we have limit spaces $X(G_0)$ and $X(G)$. A homeomorphism $f\colon X(G_0) \to X(G)$ is called a *rearrangement* if there exist finitely many cells $C(\varepsilon_1), \dots, C(\varepsilon_n)$ such that the cells cover $X(G_0)$, have pairwise disjoint interiors, and such that each restriction $f|_{C(\varepsilon_i)}$ is a canonical homeomorphism.

The category whose objects are the $X(G)$ and whose morphisms are rearrangements is a groupoid. Note that depending on the choices of G_0 and G, a rearrangement $X(G_0) \to X(G)$ might not exist, thus the groupoid is not connected.

We can also extend the definition of graph pair diagram from Definition 1.8:

Definition 1.12 (Graph pair diagram)**.** Let $f\colon X(G_0) \to X(G)$ be a rearrangement. A *graph pair diagram* for f is a triple (E_-, E_+, φ) where E_- is an expansion of G_0, E_+ is an expansion of G and $\varphi\colon E_- \to E_+$ is a graph isomorphism, such that for every edge ε of E_- the restriction of f to $C(\varepsilon)$ is a canonical homeomorphism from $C(\varepsilon)$ to $C(\varphi(\varepsilon))$.

Rearrangements can be decomposed into certain fundamental rearrangements.

Definition 1.13 (Special rearrangements)**.** If G is an expansion of G_0, there is a canonical rearrangement $\varphi\colon X(G_0) \to X(G)$ with diagram (G, G, id). We call it an *expansion rearrangement*. Its inverse is a *contraction rearrangement*. If the expansion was simple, we also say that the expansion/contraction rearrangement is *simple*. If $\varphi\colon G_0 \to G$ is a graph isomorphism then the diagram (G_0, G, φ) represents a rearrangement $X(G_0) \to X(G)$, also denoted φ, which is called a *base isomorphism*.

Thus any rearrangement is a product of an expansion rearrangement, a base isomorphism, and a contraction rearrangement.

We are now approaching the definition of the cube complex $K(G_0, e \to R)$. The fundamental objects here are rearrangements with a fixed domain $X(G_0)$.

Definition 1.14 (Range equivalence, expansion/contraction)**.** Let $f\colon X(G_0) \to X(G_1)$ and $g\colon X(G_0) \to X(G_2)$ be two arrangements. We say that f and g are *range equivalent* if there is a base isomorphism $\varphi\colon X(G_1) \to X(G_2)$ such that $\varphi \circ f = g$. We write $[f]$ for the range equivalence class of f. We say that g is a *(simple) expansion* of f if there is a (simple) expansion rearrangement $\varphi\colon X(G_1) \to X(G_2)$ such that $\varphi \circ f = g$. In that case f is a *(simple) contraction* of g. We also apply these notions to $[f]$ and $[g]$. In special cases we will need notation for this. If $S = \{\varepsilon_1, \dots, \varepsilon_k\} \subseteq E(G_1)$ is the set of edges of G_1 such

$G_2 = G_1 \lhd \varepsilon_1 \lhd \dots \varepsilon_k$ (the order does not matter) then we write $\varphi = \Delta_S$. Conversely if $T = \{R_1, \dots, R_k\}$ are the replacements for the ε_i then we write $\varphi^{-1} = \nabla_T$. If such an S or T has only one element, we may omit the set braces from the notation, and so write things like Δ_ε and ∇_R. In that case we may specify the subgraph R of G_2 by listing its edges, writing for example $\nabla_{\varepsilon 1, \varepsilon 2, \varepsilon 3}$.

Definition 1.15 (The cube complex)**.** Let $K(G_0, e \to R)$ be the cube complex defined as follows. There is a 0-cube for every range equivalence class $[f]$ of rearrangements with domain $X(G_0)$. For each 0-cube $[f]$, say represented by the rearrangement $f\colon X(G_0) \to X(G)$, and for each $S \subseteq E(G)$, there is an $|S|$-dimensional cube whose 0-subcube set is $\{[\Delta_T \circ f] \mid T \subseteq S\}$.

For example, if $f\colon X(G_0) \to X(G)$ is a rearrangement and ε is an edge of G then $[f]$ and $[\Delta_\varepsilon \circ f]$ span a 1-cube. Belk and Forrest proved that $K(G_0, e \to R)$ is a CAT(0) cube complex:

Proposition 1.16 ([BF15b, Proposition 3.33, Corollary 3.24])**.** *The complex $K(G_0, e \to R)$ is contractible, and in fact is* CAT(0).

Definition 1.17 (Rank)**.** The *rank* of a rearrangement $f\colon X(G_0) \to X(G)$ is $\mu(f) := |E(G)|$. Since range equivalent rearrangements have the same rank, we can also define the rank of an equivalence class $\mu([f]) := \mu(f)$.

It is easy to see (and explained in [BF15b]) that μ extends to a Morse function on the cube complex $K(G_0, e \to R)$ (for any graph G_0 and replacement system $e \to R$). We review the relevant Morse theoretic concepts in the next subsection.

1.4. **Morse theory.**

Definition 1.18 (Morse function)**.** Let X be a cube complex and let $h\colon X \to \mathbb{R}$ be a map. We call h a *Morse function* if the following properties hold:

(i) The image $h(X^{(0)})$ is discrete in $\mathbb{R}$.
(ii) For any cube c, the restriction of h to c is an affine function $h|_c\colon c \to \mathbb{R}$.
(iii) For any cube c of positive dimension, the restriction of h to c is non-constant.

Here when we say that $h|_c\colon c \to \mathbb{R}$ is an affine function, we are viewing c as $[-1, 1]^{\dim(c)}$ with the usual affine structure.

Given a cube complex X with a Morse function h, for any $m \in \mathbb{R}$ we denote by X_m the subcomplex of X supported on those 0-cubes x with $h(x) \leq m$. The *descending link* $\mathrm{lk}{\downarrow}x$ of a 0-cube x is the link of x in the sublevel set $X_{h(x)}$. That is, the descending link is the subcomplex of the link $\mathrm{lk}\, x$ supported on those 0-simplices along which h is decreasing. Our Morse Lemma (which is a special case of [BB97, Corollary 2.6]) is as follows:

Lemma 1.19 (Morse Lemma)**.** *Let X be a cube complex and $h\colon X \to \mathbb{R}$ a Morse function. Suppose $m \leq M$ are real numbers, and that for all 0-cubes $x \in X^{(0)}$ with $m < h(x) \leq M$*

the descending link $\mathrm{lk}{\downarrow}x$ *is* $(n-1)$*-acyclic. Then the inclusion* $X_m \to X_M$ *induces an isomorphism in* H_k *for* $k \leq n-1$*, and an epimorphism in* H_n*.*

The connection between Morse theory and finite presentability is made as follows. First recall that if a group is finitely presented then it satisfies the homological finiteness property of being of type FP_2. To show that T_B is not of type FP_2 we will use the following criterion.

Lemma 1.20. *Let a group* Γ *act with finite stabilizers on a* 1*-acyclic cube complex* X*. Let* $h\colon X \to \mathbb{R}$ *be a Morse function on* X*, and suppose that each* X_m *is* Γ*-invariant and cocompact. Suppose there exists* $N \in \mathbb{N}$ *such that for any* $x \in X^{(0)}$ *with* $h(x) > N$*, the descending link of* x *is connected. Then* Γ *is of type* FP_2 *if and only if* X_m *is* 1*-acyclic for some* $m \geq N$*.*

Proof. If some X_m is 1-acyclic then it is of type FP_2 by [Bro87, Proposition 1.1]. In the converse direction, the Morse Lemma implies that the maps $H_1(X_N \to X_m)$ are all surjective for $m \geq N$. Thus if X_m is not 1-acyclic for any m, then the map $H_1(X_n \to X_m)$ cannot be trivial for any $N \leq n \leq m$ (since the previous map is surjective and factors through this one). Hence the system $H_1(X_n), n \geq N$ is not essentially trivial, and [Bro87, Theorem 2.2] implies that G is not of type FP_2. □

We close the section with a strengthening of the well known nerve lemma. We will only need the case $n = 1$, where it is a homological version of [WZ15, Lemma 6.2], but having the question from the introduction in mind, we prove the general case. The main point is the surjective morphism at the end of the statement.

Proposition 1.21 (Strong nerve lemma)**.** *Let a CW complex* X *be covered by subcomplexes* $(X_i)_{i\in I}$*, let* L *be the nerve of the cover and let* $n \in \mathbb{N}$ *be arbitrary. Assume that if* $X_{i_1} \cap \ldots \cap X_{i_r} \neq \emptyset$ *then*

$$H_q(X_{i_1} \cap \ldots \cap X_{i_r}) = H_q(pt.)$$

for $1 \leq r \leq n-q$*. Then* $H_k(X) \cong H_k(L)$ *for* $k < n$ *and there is a surjective morphism* $H_n(X) \to H_n(L)$*.*

Proof. The nerve is the simplicial complex of those $T \subseteq I$ with $X_T := \bigcap_{i\in T} X_i \neq \emptyset$. It is equipped with the coefficient systems $\mathfrak{h}_q\colon T \to H_q(X_T)$. We use the spectral sequence from [Dav08, Theorem E.3.2]:

$$E^2_{p,q} = H_p(L; \mathfrak{h}_q) \Rightarrow H_*(X).$$

Our conditions ensure that $\mathfrak{h}_q(T) = 0$ provided T has dimension at most $n-q-1$ for $q > 0$. This means that

$$E^2_{p,q} = \begin{cases} H_p(L; \mathfrak{h}_0) & q = 0 \text{ and } 0 \leq p \leq n \\ 0 & 1 \leq q \leq n-p-1. \end{cases} \tag{1.1}$$

Since the spectral sequence lies in the first quadrant, the region $p + q \leq n-1$ of (1.1) remains stable, which tells us that $H_p(X) \cong H_p(L; \mathfrak{h}_0)$ for $0 \leq p \leq n-1$ and that there is an epimorphism $H_n(X) \to H_n(L; \mathfrak{h}_0)$.

Thus it remains to show that $H_p(L;\mathfrak{h}_0) \cong H_p(L)$ for $0 \leq p \leq n-1$ and to produce an epimorphism $H_n(L;\mathfrak{h}_0) \to H_n(L)$. We claim that the map of coefficient systems $\mathfrak{h}_0 \to \mathbb{Z}$ (here the codomain is the coefficient system that is constant $\mathbb{Z}$) induced by contracting all X_T to a point works. To spell this out let $(C_p)_{p\in\mathbb{N}}$ be the chain complex of L, i.e. $C_p \cong \bigoplus_{T\in\Sigma_p} \mathbb{Z}$ where Σ_p is the set of p-simplices of L. We think of the generator of the summand corresponding to T as X_T. Let $\tilde{C}_p$ be the chain complex of L with coefficients in $\mathfrak{h}_0$, i.e. $\tilde{C}_p \cong \bigoplus_{T\in\Sigma_p} H_0(X_T)$. Our assumptions ensure that $H_0(X_T) = \mathbb{Z}$ for T of dimension at most $n-1$, so we can identify C_p with $\tilde{C}_p$ for $p \leq n-1$. It follows that $H_p(L) \cong H_p(L;\mathfrak{h}_0)$ for $p \leq n-2$. It is also easy to see that the images of $C_n \xrightarrow{\partial} C_{n-1}$ and of $\tilde{C}_n \xrightarrow{\partial} C_{n-1}$ are the same so $H_{n-1}(L) \cong H_{n-1}(L;\mathfrak{h}_0)$.

In general (to cover degree n), we look at the map of chain complexes $\psi_*\colon \tilde{C}_* \to C_*$ taking each component of X_T to X_T. The key point is that this is actually a map of chain complexes, in particular that the square

$$\begin{array}{ccc} \tilde{C}_{n+1} & \xrightarrow{\partial} & \tilde{C}_n \\ \downarrow{\scriptstyle \psi_{n+1}} & & \downarrow{\scriptstyle \psi_n} \\ C_{n+1} & \xrightarrow{\partial} & C_n \end{array}$$

commutes. This comes from the fact that if $f\colon X_U \to X_T$ is continuous (inclusion in our case) then the diagram

$$\begin{array}{ccc} H_0(X_U) & \xrightarrow{H_0(f)} & H_0(X_T) \\ \downarrow & & \downarrow \\ H_0(\text{pt.}) & \longrightarrow & H_0(\text{pt.}) \end{array}$$

commutes. The rest is easy: any $c \in C_n$ is the image under ψ_n of some $\tilde{c} \in \tilde{C}_n$ and if $\tilde{c} = \partial\tilde{d}$ then $c = \partial\psi_{n+1}(\tilde{d}) =: \partial d$. Thus ψ_* induces an epimorphism in degree n. □

Corollary 1.22. *Let a CW complex X be covered by connected subcomplexes. If the nerve of the cover is not 1-acyclic then neither is X.* □

1.5. CAT(0) **cube complexes.** In this brief subsection we collect some terminology and results regarding CAT(0) cube complexes, all of which comes for example from [HW08].

Let X be a CAT(0) cube complex. A *midcube* is the subset of $[-1,1]^n$ obtained by restricting some coordinate to 0. The midcubes of X form a new cube complex and each of its components is a *hyperplane* in X. Each hyperplane is naturally embedded in X. A *wall* of X is a parallel class of oriented 1-cubes of X. We will also use the term *wall* to denote

the cube complex whose 0-cubes are those 1-cubes, whose 1-cubes are the corresponding 2-cubes of X, and so forth. There is a natural 2-to-1 correspondence between walls and hyperplanes and each wall is isomorphic to its corresponding hyperplane. If c is a cube and H is a wall containing a 1-face of c we say that H *cuts through* c. The 1-faces of a k-cube lie in precisely k walls. Given a hyperplane H, any subcomplex of X supported on all the 0-cubes in a connected component of $X \setminus H$ is called a *half-space*. In fact each wall H determines two half-spaces and we denote them H^+ and H^-.

2. Proof

In this section we prove our main result, that the Basilica Thompson group T_B is not finitely presented. The main technical lemma that is specific to the Basilica rewriting system is Lemma 2.4. The Lemmas 2.6 and 2.7 should readily generalize to other setups. The rest of the section is general combinatorial topology which is conveniently phrased in the setup of cube complexes.

Remark 2.1. We think that it would only take minor modifications (mostly to Lemma 2.4) to prove that, for example, the rearrangement groups for the rabbit Julia set and its variants are not finitely presented (see [BF15b] for background). Thus a positive answer to the question posed initially will not come from Julia sets whose parameter is in a bulb of the Mandelbrot set adjacent to the main cardioid.

Using Lemma 1.20 our main task is to show that certain sublevel sets are not 1-acyclic. For that purpose we will use the following criterion.

Lemma 2.2. *Let X be a* CAT(0) *cube complex, let Y be a subcomplex, and let H_1, H_2 be two walls in X. Assume that $H_1^\delta \cap Y$ and $H_2^\delta \cap Y$ are connected for $\delta \in \{+,-\}$, that $H_1^{\delta_1} \cap H_2^{\delta_2} \cap Y$ are all non-empty for $\delta_1, \delta_2 \in \{+,-\}$, and that $H_1 \cap H_2 \cap Y$ is empty. Then Y is not* 1*-acyclic.*

Proof. The assumption that $H_1 \cap H_2 \cap Y$ is empty means that Y contains no cube through which H_1 and H_2 both cut, which implies that Y is covered by the subcomplexes $H_i^\delta \cap Y$, for $i \in \{1,2\}$ and $\delta \in \{+,-\}$. The nerve of this cover is $S^0 * S^0 \cong S^1$, by the assumption that the $H_1^{\delta_1} \cap H_2^{\delta_2} \cap Y$ are all non-empty. Since we are also assuming that each $H_i^\delta \cap Y$ is connected, and since the nerve is not 1-acyclic, the result follows from Corollary 1.22. □

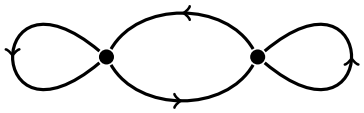

Figure 1. The graph G_0.

Now consider the cube complex $K(G_0, e \to R)$, for G_0 the graph in Figure 1 and $e \to R$ the Basilica replacement rule from Example 1.9. For the rest of this section, we will always be using this rewriting rule, so we omit it from the notation. Let $h\colon K(G_0) \to \mathbb{R}$ be the Morse function induced by the rank. To prove that T_B is not of type FP_2, Lemma 1.20 says it suffices to show that for m large enough, no $K(G_0)_m$ is 1-acylic. We will apply Lemma 2.2, namely, we will find two walls in $K(G_0)$ satisfying the requirements for $Y = K(G_0)_m$.

First we claim that walls in $K(G_0)$ are isomorphic to cube complexes of the form $K(G)$, for appropriate G.

Lemma 2.3 (Modeling walls)**.** *Let G_1 be a graph, let ε an edge of G_1, and let $G_2 = G_1 \lhd \varepsilon$. Let $f\colon X(G_0) \to X(G_1)$ and $g = \Delta_\varepsilon \circ f\colon X(G_0) \to X(G_2)$ be rearrangements, and let c be the 1-cube spanned by $[f]$ and $[g]$. Then the wall of c is isomorphic to the complex $K(G_1 \setminus \varepsilon)$.*

Proof. Let H be the wall of c in $K(G_0)$, so the 0-cubes of H are the 1-cubes of $K(G_0)$ parallel to c. Let d be a such a 1-cube of H. Say d has 0-faces $[f']$ and $[g']$, with $g' = \Delta_\varepsilon \circ f'$. Since d is parallel to c, we know that $f' = \phi \circ f$ for some rearrangement $\phi\colon X(G_1) \to X(G')$, and that ϕ is a composition of expansions and contractions that never involve the edge ε.

Now, with all the above data for 0-cubes d in H, define a map $\Psi\colon H^{(0)} \to K(G_1 \setminus \varepsilon)^{(0)}$ by

$$\Psi(d) := [\phi|_{K(G_1 \setminus \varepsilon)}].$$

Here $\phi|_{K(G_1\setminus\varepsilon)}$ is, as the notation implies, the restriction of ϕ to a rearrangement $X(G_1 \setminus \varepsilon) \to X(G' \setminus \varepsilon)$. This is well defined since ϕ is a composition of expansions and contractions that never involve ε. It is easily seen that Ψ is bijective, and that it extends to a cubical isomorphism $\Psi\colon H \to K(G_1 \setminus \varepsilon)$. □

Our specific walls of interest arise as follows. Consider the family of graphs J_n with $n + 5$ vertices and $2n + 10$ edges indicated in Figure 2.

Fix a vertex $[f_n]$ of $K(G_0)$ represented by a rearrangement $X(G_0) \to X(J_n)$. These exist because J_n is an expansion of G_0. Let c_n be the 2-cube whose 0-faces are $[f_n]$, $[\nabla_{c,d,e} \circ f_n]$, $[\nabla_{x,y,z} \circ f_n]$ and $[\nabla_{c,d,e} \circ \nabla_{x,y,z} \circ f_n]$. Let H_1 be the hyperplane that cuts through the 1-cube from $[f_n]$ to $[\nabla_{x,y,z} \circ f_n]$ and H_2 the hyperplane that cuts through the 1-cube from x_n to $[\nabla_{c,d,e} \circ f_n]$. For $i = 1, 2$, let H_i^+ be the half-spaces containing $[f_n]$ and let H_i^- be the half-spaces that do not contain $[f_n]$.

By Lemma 2.3, the intersection $H_1 \cap H_2$ is isomorphic to $K(O_n)$, where O_n is the graph given in Figure 3.

The following is our key technical lemma, and is the result that is most specific to the case of the Basilica rewriting system.

Lemma 2.4 (Key technical lemma)**.** *Let $n \in \mathbb{N}$ be arbitrary. Let O_n be the graph on $n+3$ vertices and $2n + 4$ edges indicated in Figure 3. Then the sublevel complex $K(O_n)_{2n+3}$ is empty.*

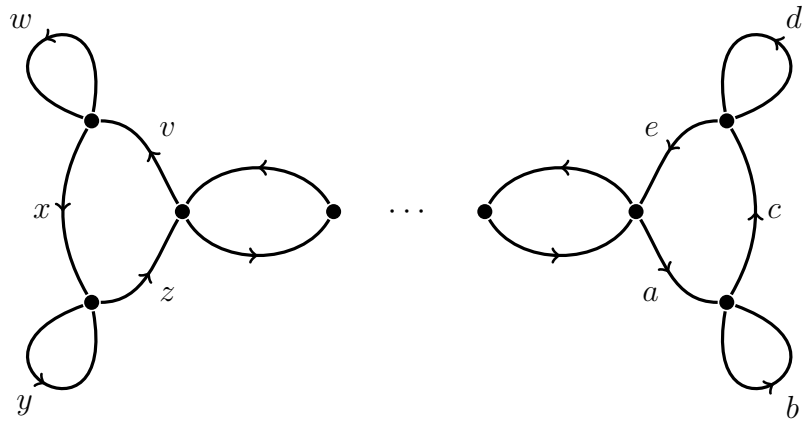

FIGURE 2. The family of graphs $J_n, n \in \mathbb{N}$ with various edges labeled. The graph J_n has $n+5$ vertices and $2n+10$ edges.

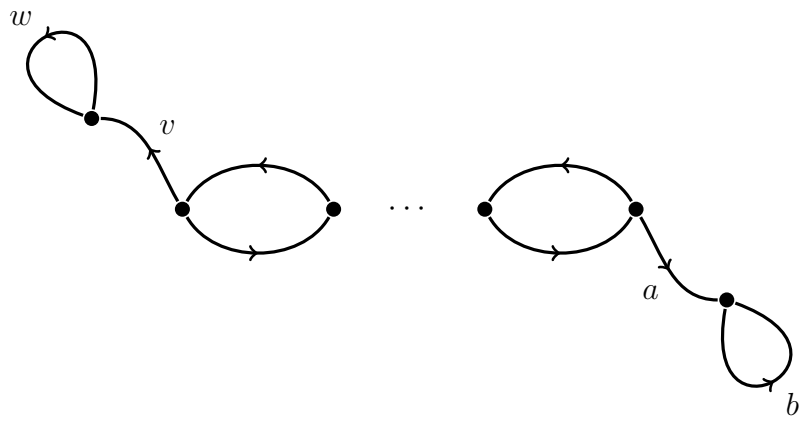

FIGURE 3. The family of graphs $O_n, n \in \mathbb{N}$. The graph O_n has $n+3$ vertices and $2n+4$ edges.

Proof. We need to show that any graph obtained from O_n via a sequence of expansions and contractions has at least $2n+4$ edges. Call such a graph *relevant*. Note first that for any relevant graph, denoting by V the number of vertices and E the number of edges, we have $E-(2n+4) = 2(V-(n+3))$. Thus we must equivalently show that any relevant graph has at least $n+3$ vertices. Note further that in a relevant graph all edges lie in the boundary of the outer region, so they inherit a cyclic ordering. We call a path $e_1, \ldots, e_k$ in a relevant graph *ordered* if the edges are successive in the cyclic order.

We say that a non-constant ordered closed path in a relevant graph G is a *special circuit* if it only meets vertices of degree 4. We call a vertex v of such a graph *collapsible* if it is of degree 4 and there is a special circuit connecting v to itself. We see that O_n has zero collapsible vertices. Intuitively, if v is a collapsible vertex then after repeated contractions it can be made to disappear.

Let G be a relevant graph, let ε be an edge of G and let $G \lhd \varepsilon$ be the corresponding simple expansion. Then the new vertex of $G \lhd \varepsilon$ is collapsible: the special circuit in question is just the new loop, $\varepsilon 2$. Any vertex of G that is collapsible is still collapsible in $G \lhd \varepsilon$: a special circuit in G that uses ε can be turned into a special circuit in $G \lhd e$ by replacing ε by $\varepsilon 1, \varepsilon 2, \varepsilon 3$. Finally, if a vertex v of G is not collapsible, then it does not become collapsible in $G \lhd \varepsilon$: indeed, a special circuit in $G \lhd \varepsilon$ at v that uses one of $\varepsilon 1$, $\varepsilon 2$ or $\varepsilon 3$ has to use all of them, so replacing the sequence $\varepsilon 1, \varepsilon 2, \varepsilon 3$ by ε gives rise to a special circuit in G.

The result of this discussion is that the number of collapsible vertices minus the number of vertices is invariant under arbitrary expansions and contractions. Since the number of collapsible vertices cannot become negative, we conclude that any graph obtained from O_n via expansions and contractions has to have at least $n + 4$ vertices. □

Let $Y_n := K(G_0)_{2n+9}$. We will now verify that for n sufficiently large, all the hypotheses of Lemma 2.2 hold, with Y_n playing the role of Y and H_1 and H_2 being our two walls.

Corollary 2.5 (Empty middle)**.** *The intersection $H_1 \cap H_2 \cap Y_n$ is empty.*

Proof. By Lemma 2.3 the intersection $H_1 \cap H_2$ is isomorphic to $K(O_n)$, and $H_1 \cap H_2 \cap Y_n$ is isomorphic to $K(O_n)_{2n+3}$ since a point in $K(O_n)$ with rank r corresponds to a 2-cube of $H_1 \cap H_2$ with maximum rank $r + 6$. By Lemma 2.4 $K(O_n)_{2n+3}$ is empty. □

Lemma 2.6 (Non-empty quarter-spaces)**.** *The intersections $H_1^{\delta_1} \cap H_2^{\delta_2} \cap Y_n$ are non-empty for any $\delta_1, \delta_2 \in \{+, -\}$.*

Proof. The vertices $[\nabla_{a,b,c} \circ \nabla_{v,w,x} \circ f_n]$, $[\nabla_{c,d,e} \circ \nabla_{v,w,x} \circ f_n]$, $[\nabla_{a,b,c} \circ \nabla_{x,y,z} \circ f_n]$ and $[\nabla_{c,d,e} \circ \nabla_{x,y,z} \circ f_n]$ respectively lie in $H_1^+ \cap H_2^+ \cap Y_n$, $H_1^+ \cap H_2^- \cap Y_n$, $H_1^- \cap H_2^+ \cap Y_n$ and $H_1^- \cap H_2^- \cap Y_n$, and each of these vertices has rank $2n + 6$, so lies in Y_n. □

Lemma 2.7 (Connected half-spaces)**.** *Let $n \in \mathbb{N}$. For $i \in \{1, 2\}$ and $\delta \in \{+, -\}$, we have that $H_i^\delta \cap Y_n$ is connected.*

Proof. We know from [BF15b, Example 4.5] that the only 0-cubes $x \in H_i^\delta$ with disconnected descending link have at most 4 vertices. Thus if $h(x) > 2n + 9$ the descending link in $K(G_0)$ is connected. Now H_i^- is contractible (being convex) and for every vertex $x \in H_i^-$ the descending link in H_i^- is the same as the descending link in $K(G_0)$. So the Morse argument shows that $H_i^- \cap Y_n$ is connected.

Now we look at H_i^+. We take $i = 1$ for concreteness, but the case $i = 2$ is analogous. Let $x_1, x_2 \in (H_1^+ \cap Y_n)$ be 0-cubes. Since H_1^+ is convex and thus contractible, there is an

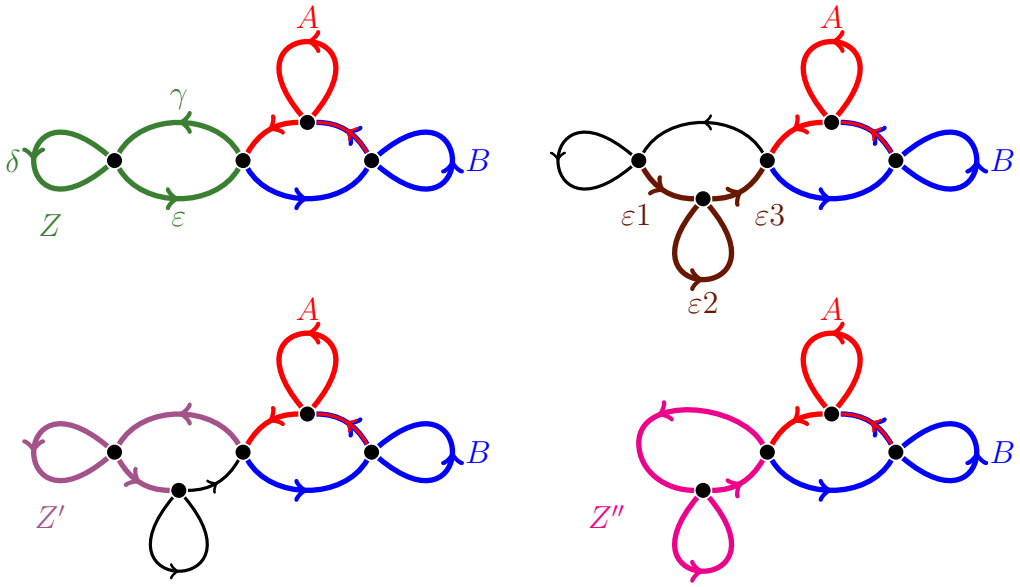

FIGURE 4. An example of the situation at the end of the proof of Lemma 2.7.

edge path from x_1 to x_2 in H_1^+. Our goal is to perturb it to lie in Y_n. For that purpose it suffices to show that given 0-cubes $x, y, z \in H_1^+$ with x adjacent to y and z, $h(x) > 2n+9$, and $h(y), h(z) < h(x)$ there exists a path from y to z using only vertices of height at most $h(x) - 1$. Let $x = [f]$ where f is a rearrangement $X(G_0) \to X(G)$ and let A and B be the (3-edge) subgraphs of G such that $y = [\nabla_A \circ f]$ and $z = [\nabla_B \circ f]$. There is a path from y to z in the descending link of x in $K(G_0)$, but we need to deal with the fact that the descending link in H_1^+ may be smaller. We only need to treat the case where this actually happens, i.e. where some 1-simplex over x lies in H_1. Say the neighbor of x in H_1^- is $x' = [\nabla_Z \circ f]$ for a subgraph Z. Moreover, if the path from y to z in the descending link does not pass through x', we are again done, so we can assume that it does pass through x' and, in fact, that y and z are both adjacent to x'. In terms of graphs this means that A and B both are edge-disjoint from Z but A and B have an edge in common.

We need some further notation: let A' be the image of A after contracting B and let B' be the image of B after contracting A. We informally describe the strategy before giving the technical details. The problem is that the path $y = [\nabla_A f], [\nabla_Z \nabla_A f], x' = [\nabla_Z f], [\nabla_Z \nabla_B f], z = [\nabla_B f]$ is not available because it does not lie in H_1^+. So rather than contracting Z we would like to first split an edge ε of Z and then contract twice (preventing us from crossing H_1). The problem with this is that $h([\Delta_\varepsilon \nabla_A f]) = h(x)$. So before we can do that we need to contract B' as well.

To make this formal, we need even more notation: say Z has edges γ, δ, ε in that order. So splitting ε in Z gives rise to a graph $Z \lhd \varepsilon$ with edges γ, δ, $\varepsilon 1$, $\varepsilon 2$ and $\varepsilon 3$ of which δ and $\varepsilon 2$ are loops. We denote by Z' the subgraph with edges γ, δ, $\varepsilon 1$ and by Z'' the result

of contracting Z' in $Z \lhd \varepsilon$. Then the vertices $[\Delta_\varepsilon f]$, $[\nabla_{Z'}\Delta_\varepsilon f]$, $[\nabla_{Z''}\nabla_{Z'}\Delta_\varepsilon f]$ all lie in H_1^+, but $h([\Delta_\varepsilon f]) = h(x) + 1$.

Finally, here is the path from y to z that does not cross H_1 and only passes through vertices of height $< h(x)$. From y it moves down to $[\nabla_{B'}\nabla_A f]$ of height $h(x)-2$. From there it moves through $[\Delta_\varepsilon \nabla_{B'}\nabla_A f]$, $[\nabla_{Z'}\Delta_\varepsilon\nabla_{B'}\nabla_A f]$ to $[\nabla_{Z''}\nabla_{Z'}\Delta_\varepsilon\nabla_{B'}\nabla_A f] = [\nabla_{B'}\nabla_A\nabla_{Z''}\nabla_{Z'}\Delta_\varepsilon f]$ of height $h(x)-3$. At this point it can "untwist" A and B by moving through $[\nabla_A\nabla_{Z''}\nabla_{Z'}\Delta_\varepsilon f]$ to $[\nabla_{Z''}\nabla_{Z'}\Delta_\varepsilon f]$ (of height $h(x)-1$) and on to $[\nabla_B\nabla_{Z''}\nabla_{Z'}\Delta_\varepsilon f]$ and $[\nabla_{A'}\nabla_B\nabla_{Z''}\nabla_{Z'}\Delta_\varepsilon f] = [\nabla_{Z''}\nabla_{Z'}\Delta_\varepsilon\nabla_{A'}\nabla_B f]$. Now proceeding symmetrically it passes through $[\nabla_{Z'}\Delta_\varepsilon\nabla_{A'}\nabla_B f]$, $[\Delta_\varepsilon\nabla_{A'}\nabla_B f]$, $[\nabla_{A'}\nabla_B f]$ to $z = [\nabla_B f]$. □

We can now prove a strengthening of our main result.

Theorem 2.8. *T_B is not of type FP_2, and hence is not finitely presented.*

Proof. Using Lemma 1.20 it suffices to show that Y_n is not 1-acyclic for arbitrarily large n. This follows from Lemma 2.2 the hypotheses of which have been verified in Corollary 2.5 and Lemmas 2.6 and 2.7. □

References

[BB97] M. Bestvina and N. Brady, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997), no. 3, 445–470.

[BF15a] J. Belk and B. Forrest, *A Thompson group for the basilica*, Groups Geom. Dyn. **9** (2015), no. 4, 975–1000.

[BF15b] ———, *Rearrangement Groups of Fractals*, arXiv:1510.03133v1, 2015.

[Bro87] K. S. Brown, *Finiteness properties of groups*, J. Pure Appl. Algebra **44** (1987), 45–75.

[CR09] P.-E. Caprace and B. Rémy, *Simplicity and superrigidity of twin building lattices*, Invent. Math. **176** (2009), no. 1, 169–221.

[Dav08] M. W. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008.

[HW08] F. Haglund and D. T. Wise, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), no. 5, 1551–1620.

[WZ15] S. Witzel and M. C. B. Zaremsky, *The Σ-invariants of Thompson's group F via Morse theory*, Topological Methods in Group Theory, London Math. Soc. Lecture Note Ser., Cambridge Univ. Press, Cambridge, 2015, To appear. arXiv:1501.06682.

Department of Mathematics, Bielefeld University, PO Box 100131, 33501 Bielefeld, Germany

E-mail address: switzel@math.uni-bielefeld.de

Department of Mathematical Sciences, Binghamton University, Binghamton, NY 13902

E-mail address: zaremsky@math.binghamton.edu